EXHIBIT B

Infringement Chart for US 6,549,988 vs. Google

Claim 1

Claim Language	Google Evidence
1. A computer suitable for use in a data storage system comprising a network interconnecting a plurality of such computers, the computer comprising	The Cloud Storage include is a computer(On-premises) that suitable for use in a data storage system (cloud storage), the data storage system comprising a network interconnecting a plurality of such computers 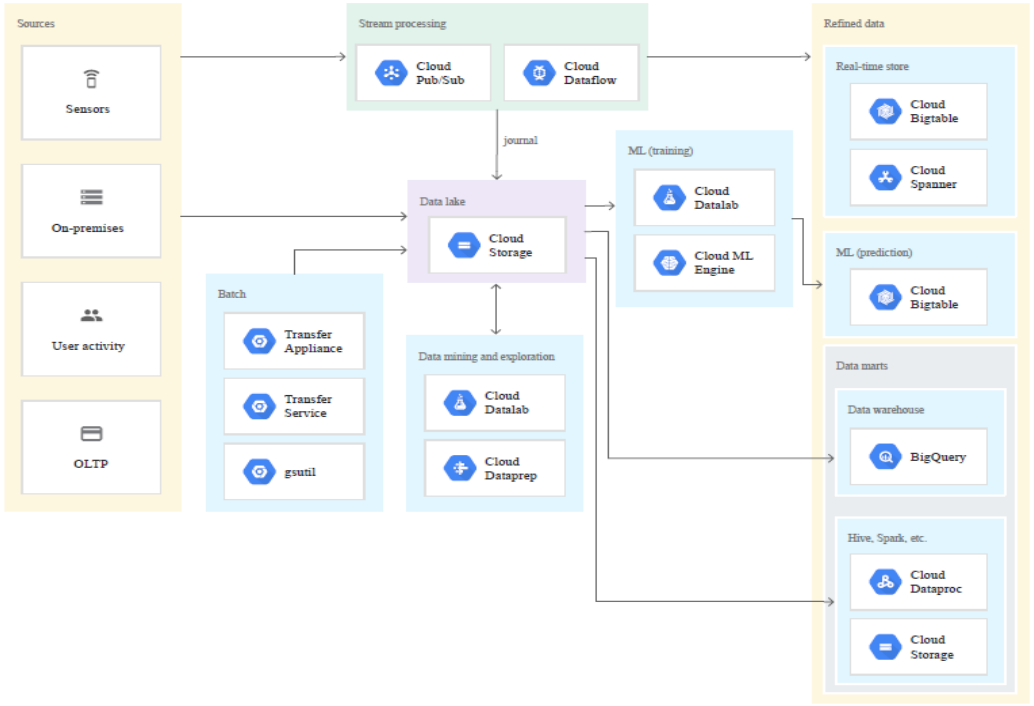

	Cloud Storage as a data lake | Cloud Architecture Center | Google Cloud Source: https://cloud.google.com/architecture/build-a-data-lake-on-gcp
an I/O channel adapter for accepting an incoming I/O request from a host;	Cloud Storage has a first interface configured to receive input/output (I/O) traffic from a first host device via a dedicated I/O channel (traffic from a host device (on-premises sources), the I/O traffic comprising a read command

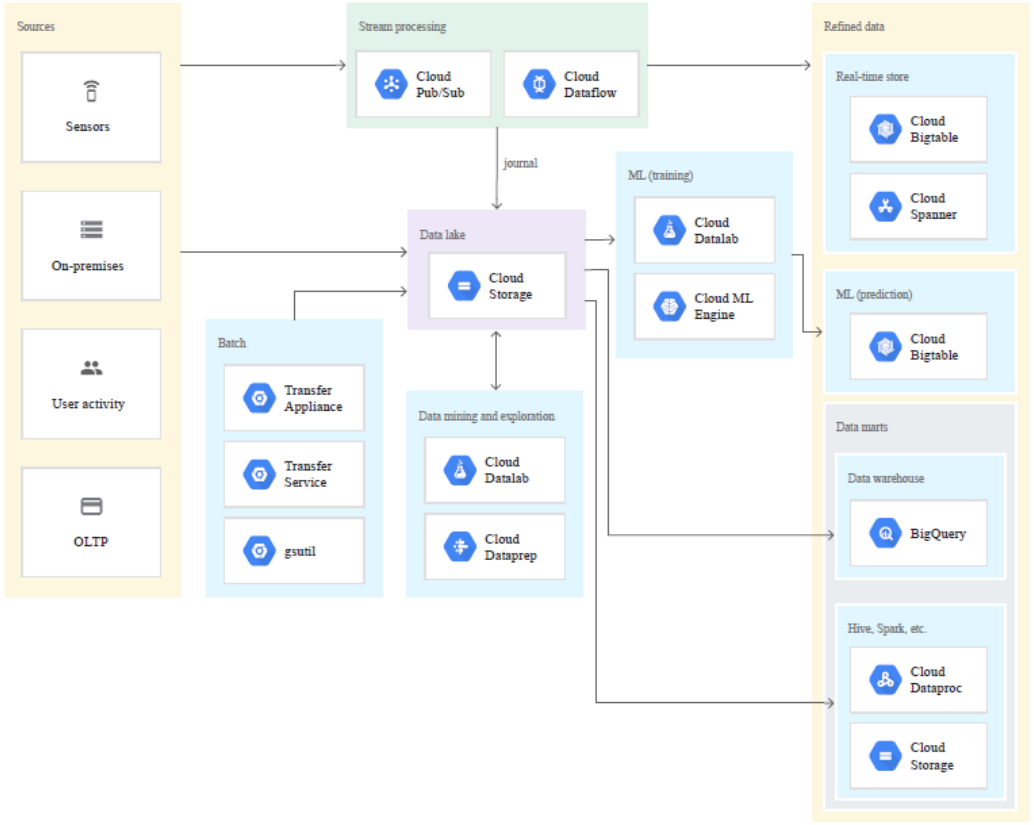

Cloud Storage as a data lake | Cloud Architecture Center | Google Cloud
Source: https://cloud.google.com/architecture/build-a-data-lake-on-gcp

configuration manager software for enabling said I/O channel adapter to decide whether (i) to route said request to cache, (ii) to route said request to disk, or (iii) to reject said request;	configuration manager software for enabling said I/O channel adapter to decide whether (i) to route said request to cache, (ii) to route said request to disk, accessed directly (**Access operation**) 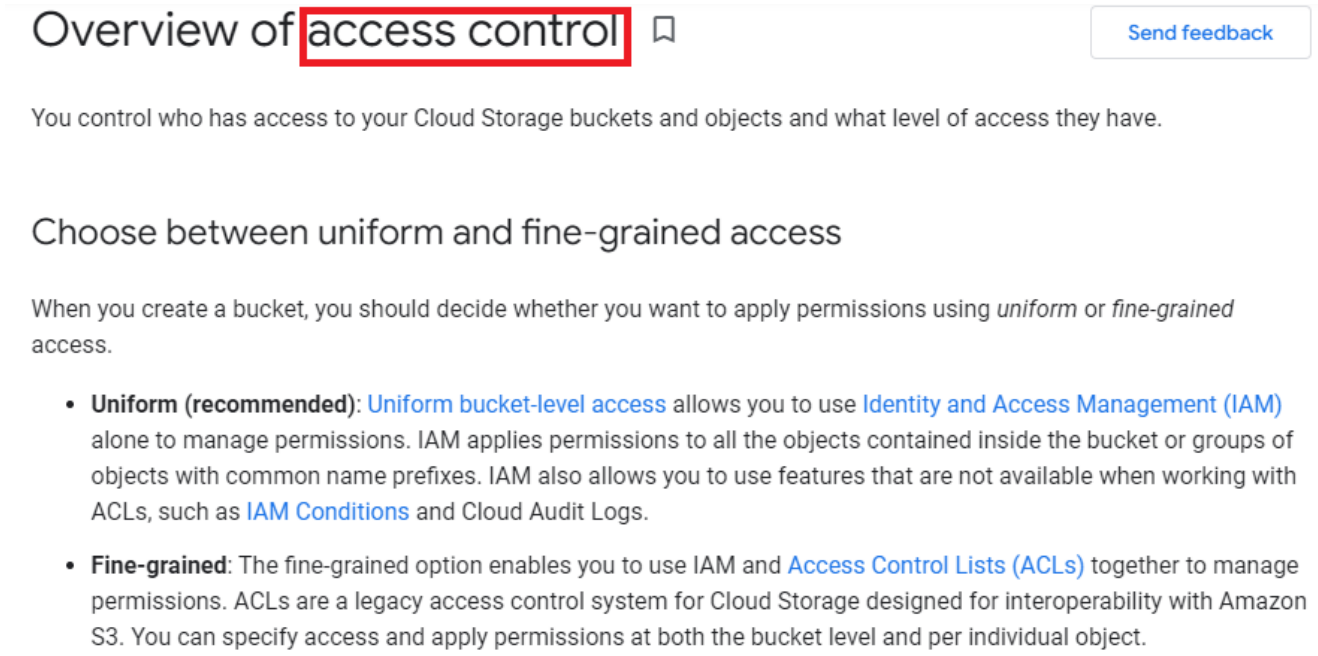 Overview of access control Send feedback You control who has access to your Cloud Storage buckets and objects and what level of access they have. Choose between uniform and fine-grained access When you create a bucket, you should decide whether you want to apply permissions using *uniform* or *fine-grained* access. • **Uniform (recommended)**: Uniform bucket-level access allows you to use Identity and Access Management (IAM) alone to manage permissions. IAM applies permissions to all the objects contained inside the bucket or groups of objects with common name prefixes. IAM also allows you to use features that are not available when working with ACLs, such as IAM Conditions and Cloud Audit Logs. • **Fine-grained**: The fine-grained option enables you to use IAM and Access Control Lists (ACLs) together to manage permissions. ACLs are a legacy access control system for Cloud Storage designed for interoperability with Amazon S3. You can specify access and apply permissions at both the bucket level and per individual object. Source: https://cloud.google.com/storage/docs/access-control
a network adapter for handling network control traffic;	Cloud Storage comprises network adapter for handling network control traffic

	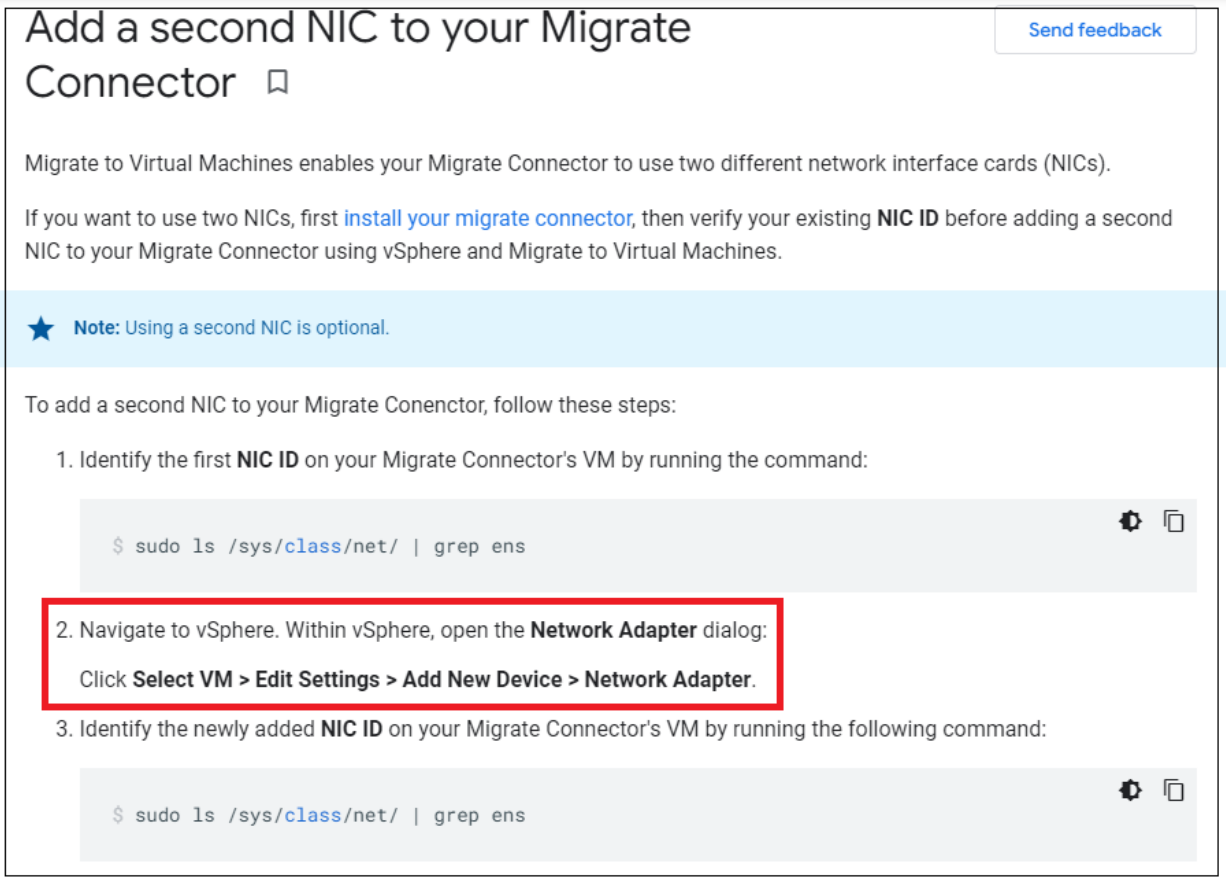 Source: https://cloud.google.com/migrate/virtual-machines/docs/5.0/how-to/second-nic
a cache memory;	Cloud Storage comprises a cache memory

	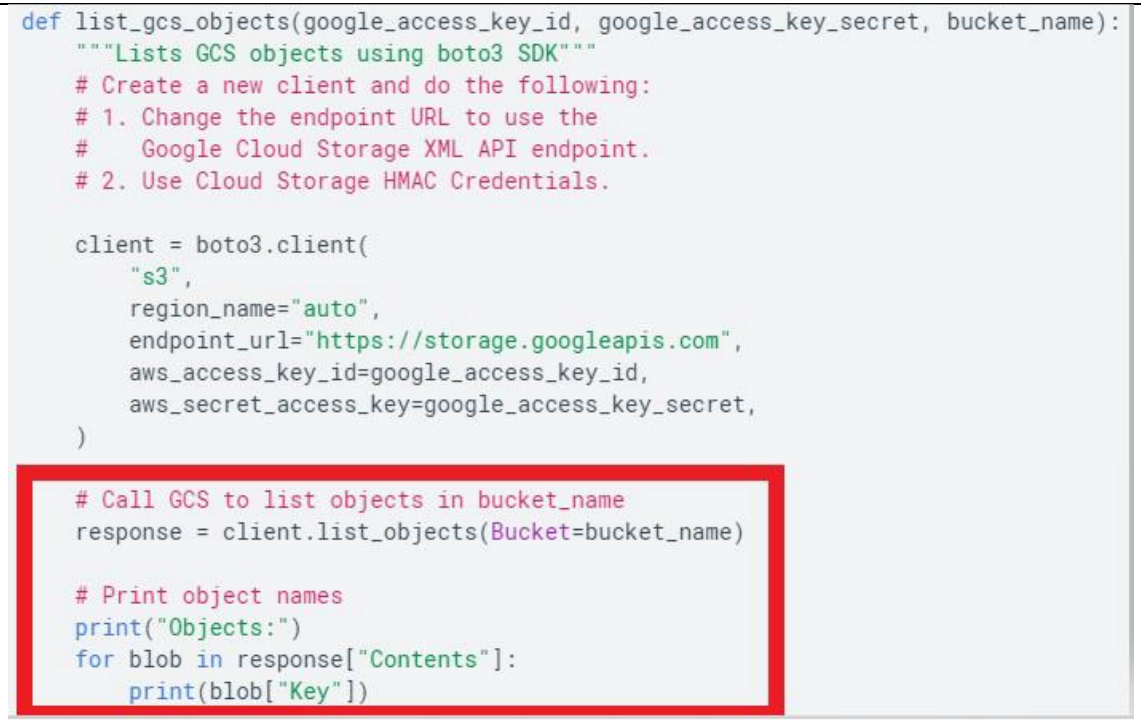 Source: https://cloud.google.com/storage/docs/migrating
front-end software for handling I/O requests arriving at the I/O channel adapter or the network adapter;	Cloud Storage comprises front-end software for handling I/O requests arriving at the I/O channel adapter or the network adapter);

```
def list_gcs_objects(google_access_key_id, google_access_key_secret, bucket_name):
    """Lists GCS objects using boto3 SDK"""
    # Create a new client and do the following:
    # 1. Change the endpoint URL to use the
    #    Google Cloud Storage XML API endpoint.
    # 2. Use Cloud Storage HMAC Credentials.

    client = boto3.client(
        "s3",
        region_name="auto",
        endpoint_url="https://storage.googleapis.com",
        aws_access_key_id=google_access_key_id,
        aws_secret_access_key=google_access_key_secret,
    )

    # Call GCS to list objects in bucket_name
    response = client.list_objects(Bucket=bucket_name)

    # Print object names
    print("Objects:")
    for blob in response["Contents"]:
        print(blob["Key"])
```

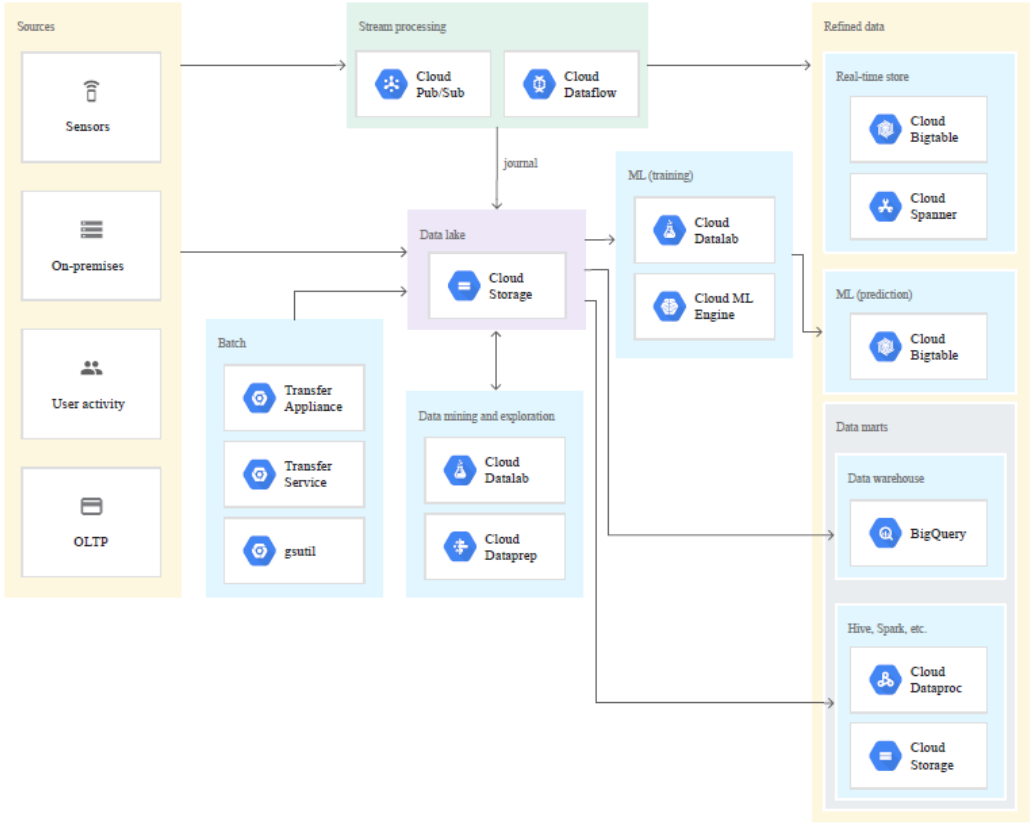

Cloud Storage as a data lake | Cloud Architecture Center | Google Cloud
Source: https://cloud.google.com/architecture/build-a-data-lake-on-gcp

cache manager software, responsive to said front-end software, for handling data stored in said cache memory; and	Cloud Storage comprises cache manager software, responsive to said front-end software, for handling data stored in said cache memory 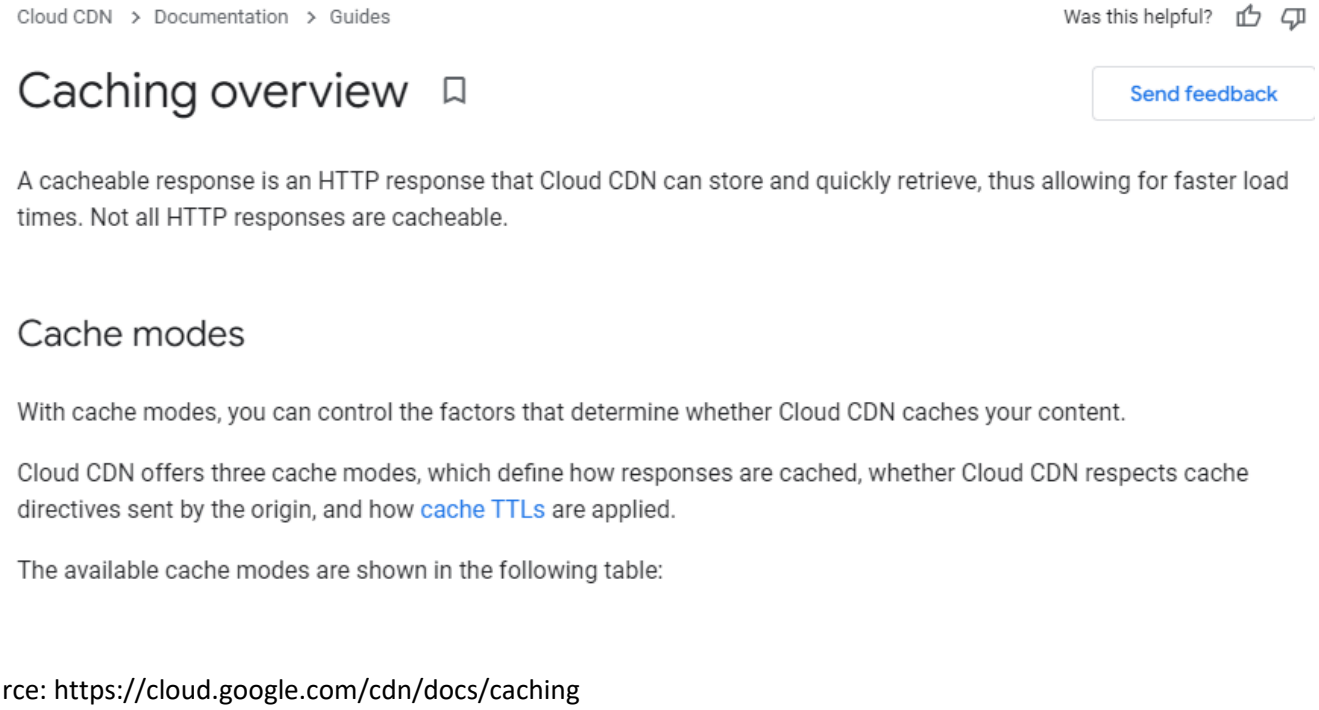 Source: https://cloud.google.com/cdn/docs/caching
back-end software, responsive to said configuration manager software, for handling reads	back-end software, responsive to said configuration manager software, for handling reads and writes to disks corresponding to the I/O requests

and writes to disks corresponding to the I/O requests but without communication over the I/O channel adapter, thereby separating disk operations from network and I/O traffic.	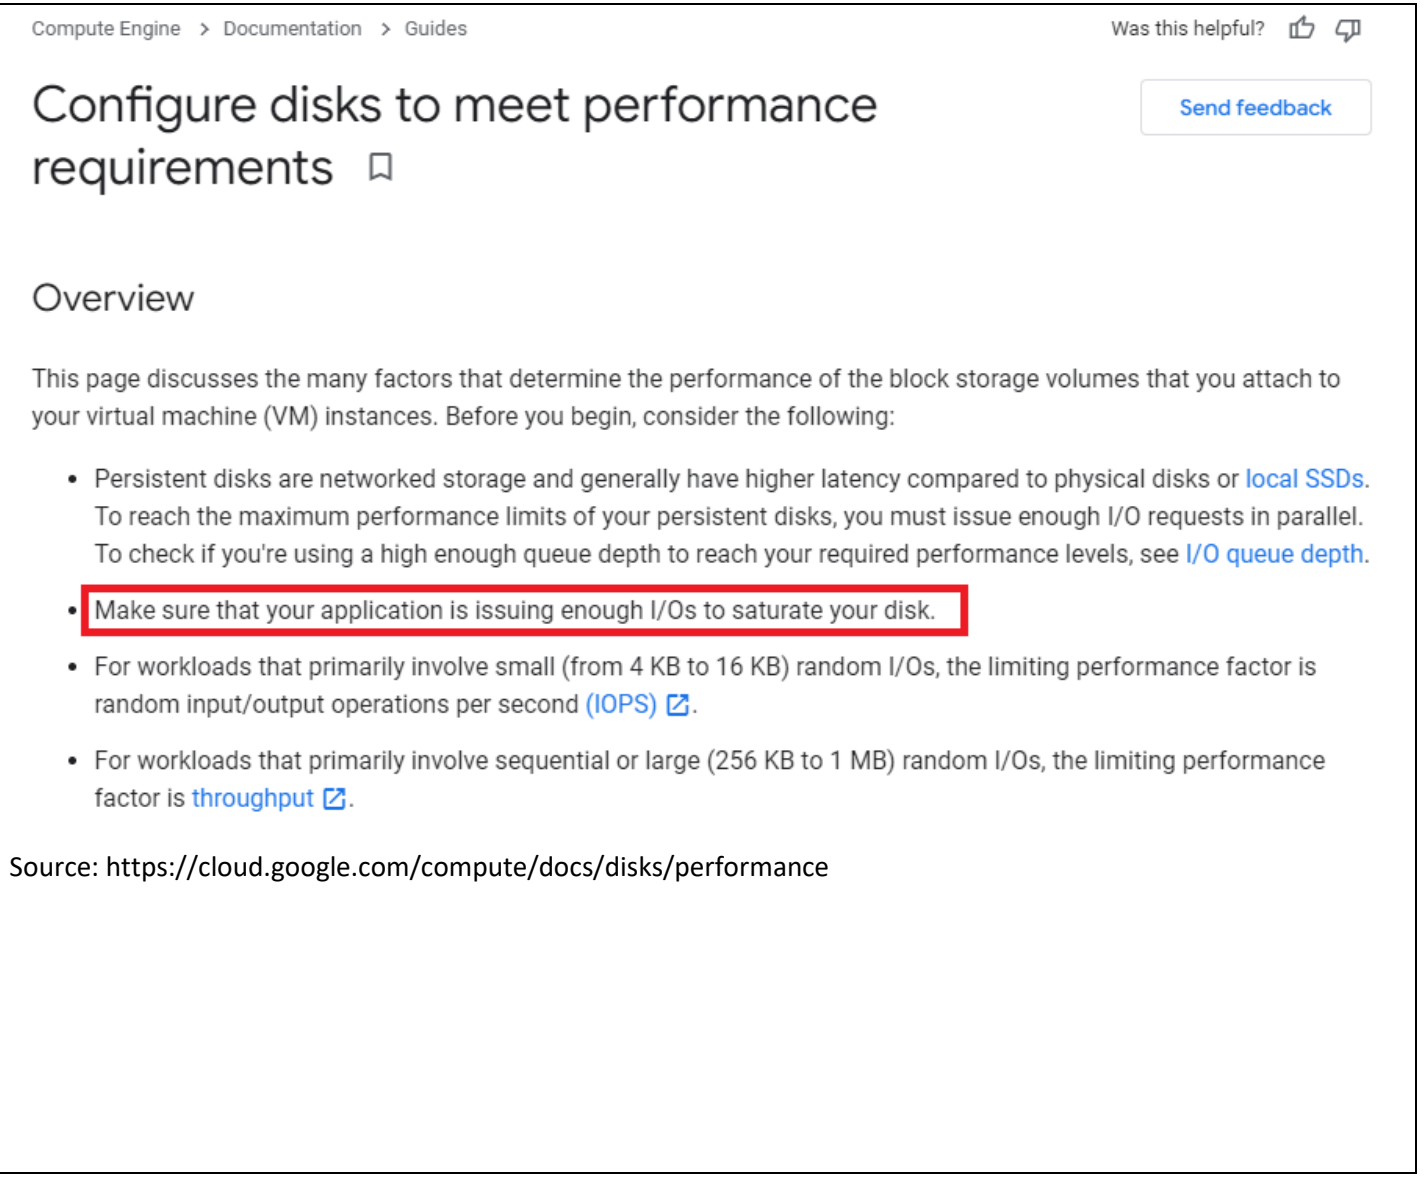 Source: https://cloud.google.com/compute/docs/disks/performance

	without communication over the I/O channel adapter, thereby separating disk operations from network and I/O traffic. 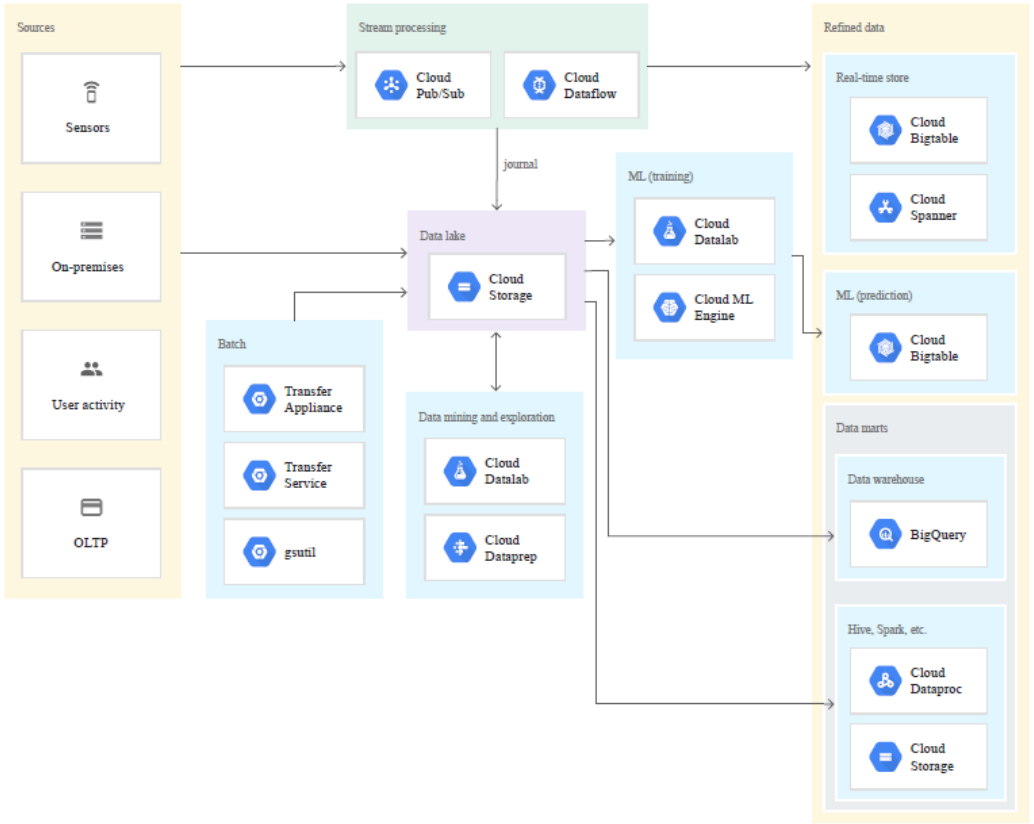 Cloud Storage as a data lake | Cloud Architecture Center | Google Cloud Source: https://cloud.google.com/architecture/build-a-data-lake-on-gcp

2. The system of claim 1 wherein the computers comprise off-the-shelf hardware and operating systems and further comprise: an adapter I/O software for accepting incoming I/O requests from a host; and a volume access table employed by the configuration manager to ensure consistency of data stored on the network.	The computer comprises off-the-shelf hardware 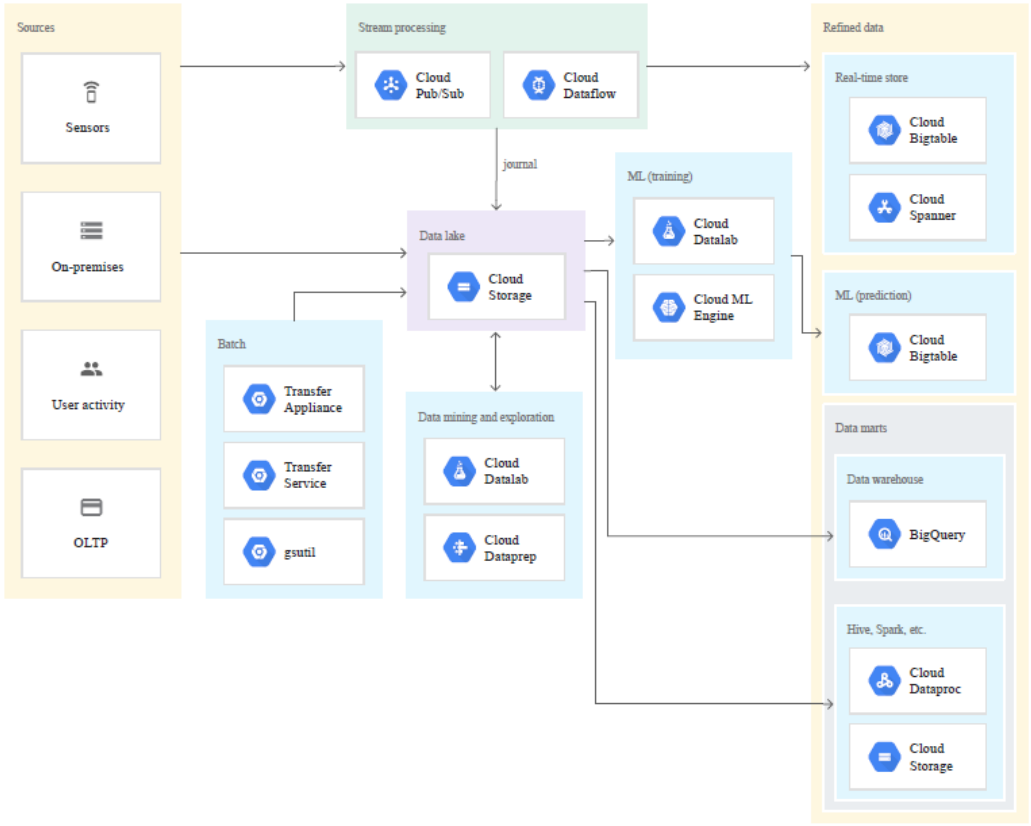 Cloud Storage as a data lake | Cloud Architecture Center | Google Cloud Source: https://cloud.google.com/architecture/build-a-data-lake-on-gcp The computers comprise off-the-shelf operating systems

	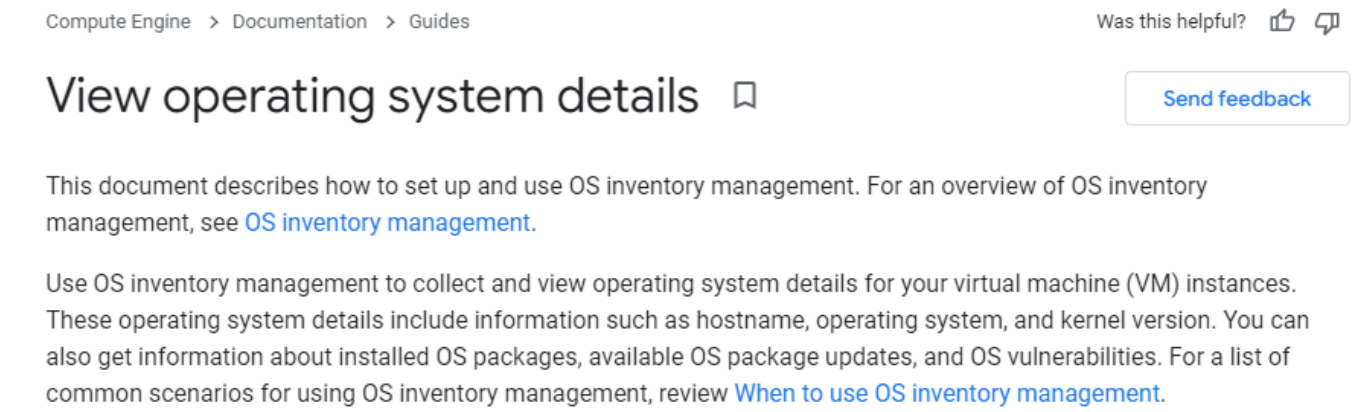 Source: https://cloud.google.com/compute/docs/instances/view-os-details
6. A method of accessing a remote disk over a computer network without incurring network overhead, the method comprising the steps of: a. causing a local host to issue a	Google Storage Transfer Service provide a method of accessing a remote disk over a computer network without incurring network overhead

request over an I/O channel to a local computer;	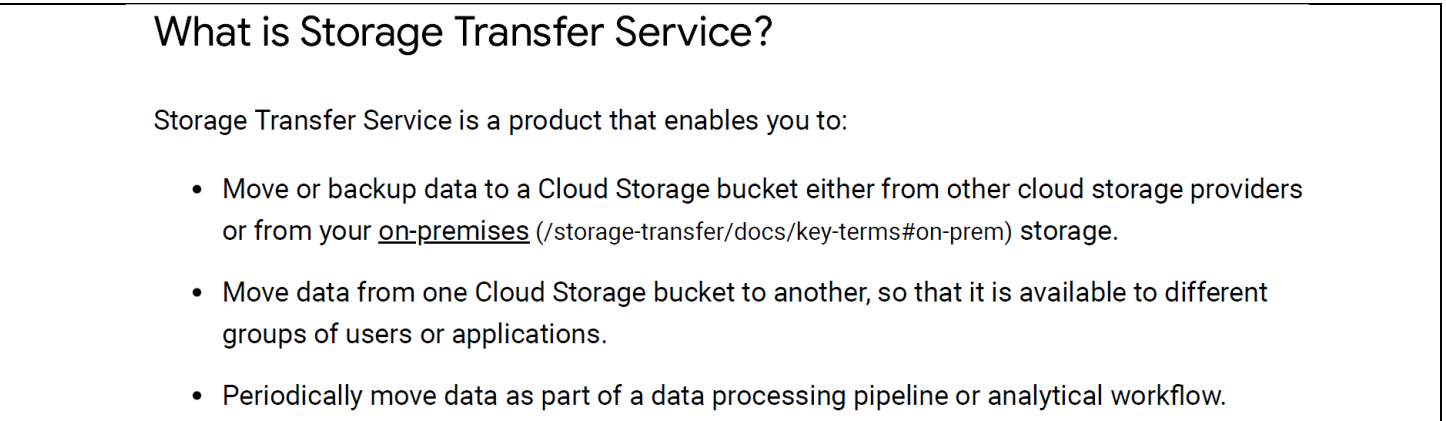 Overview | Cloud Storage Transfer Service Documentation | Google Cloud Source: https://cloud.google.com/storage-transfer/docs/overview Google Storage causes a local host to issue a request over an I/O channel to a local computer;

	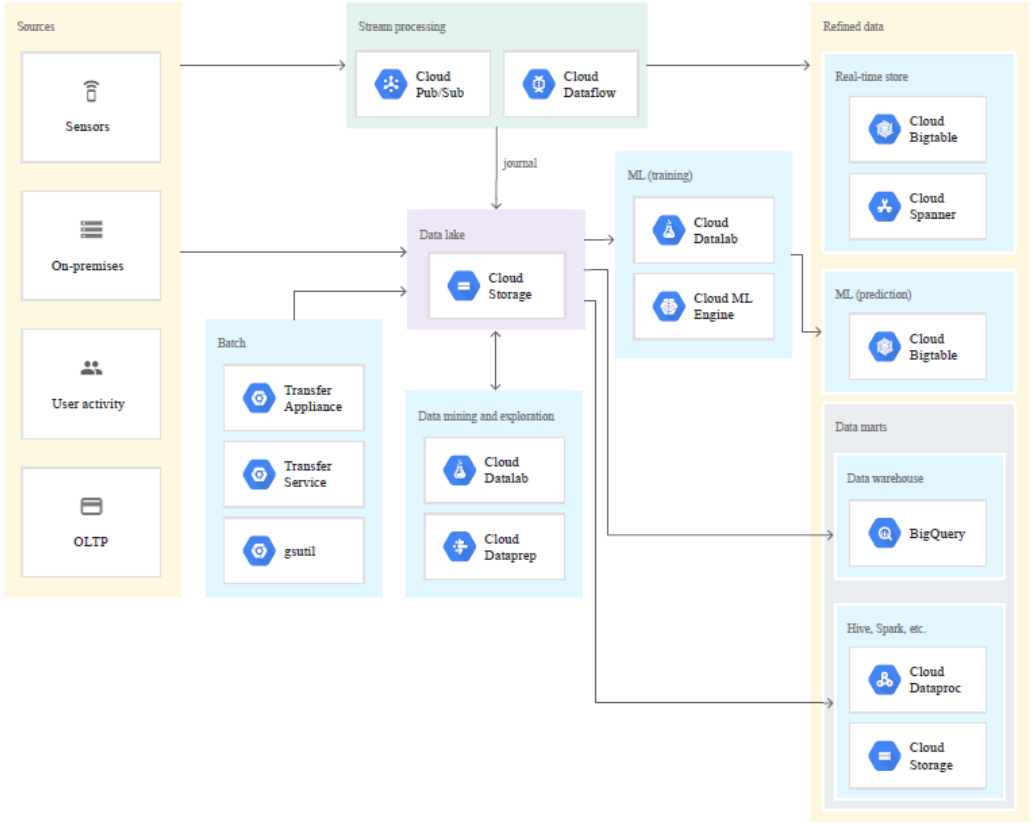 Cloud Storage as a data lake | Cloud Architecture Center | Google Cloud Source: https://cloud.google.com/architecture/build-a-data-lake-on-gcp
b. providing a configuration manager on the local computer, the	Google Storage provids a configuration manager on the local computer for Ethernet for network traffic(traffic from a host device (on-premises sources) via a dedicated I/O channel (Dedicated Interconnect Work) at a first interface).

configuration manager routing the request to a remote computer via the computer network;	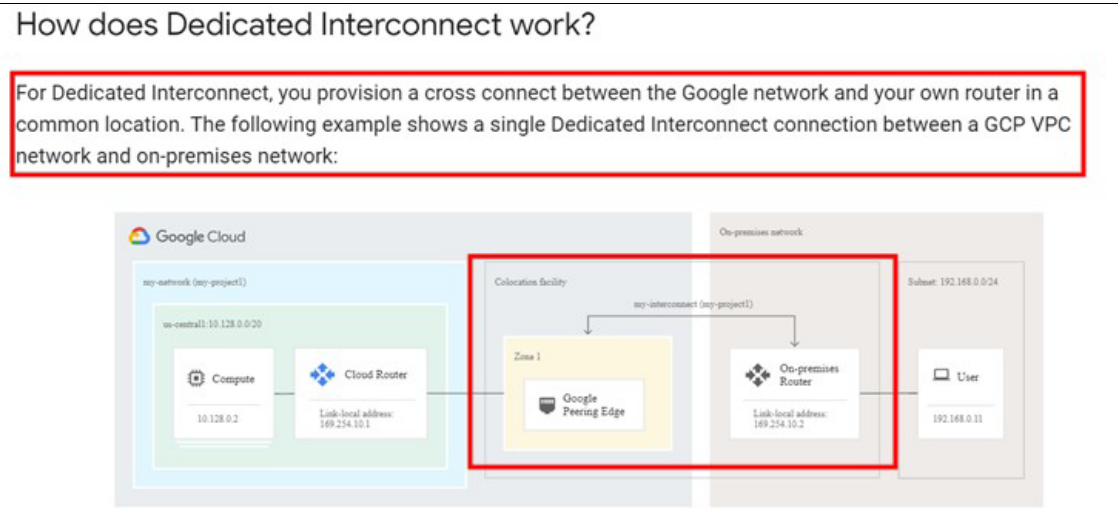 Dedicated Interconnect overview Source: https://cloud.google.com/storage-transfer/docs/on-prem-overview
c. causing the remote computer to check the request against a volume access table;	a volume access table is checked

Setting up private service access

1. Create an allocated IP address range within your VPC network for the Cloud Volumes Service mount points.

   You can't modify the IP address range after you establish it and allocate it to a volume, so we recommend allocating a range that is large enough to accommodate future usage. However, if the IP address range allocation is too small, you can add additional CIDR ranges. For more information, see Adding CIDR ranges.

   - The CVS-Performance service type needs a minimum CIDR block of `/24` (16 `/28` CIDR blocks, which each have 16 IP addresses). Some addresses in the block are used for CVS internal needs, leaving 11 addresses for your volumes from a `/28` block. Larger blocks support additional region and project pairs. For example, a `/32` block supports up to 32 combinations of region and consumer service project pairs.
   - The CVS service type (Standard-SW) needs a minimum CIDR block of `/25` (128 addresses). This supports up to the maximum 100 volumes for each project (for each zone or region, depending on the service level). A larger block can support more region and project pairs. Cross-region access isn't supported.
   - Shared VPC is supported for the CVS and CVS-Performance service types. For shared VPC networks, peering is done from the host project only. Each service project in an additional region uses an additional CIDR block of `/28` from the VPC range.

Source: https://cloud.google.com/architecture/partners/netapp-cloud-volumes/setting-up-private-services-access

d. causing the remote computer to perform an I/O operation on a disk located on the remote computer and to return data to the local computer;
e. causing the local computer to provide the returned data to the local host via the I/O channel; and
f. causing the local computer to check the data against the volume access table to ensure consistency of the data on the local and the remote computers.
	causing the remote computer to perform an I/O operation on a disk located on the remote computer (remote access) .causing the local computer to check the data against the volume access table to ensure consistency of the data on the local and the remote computers.
	 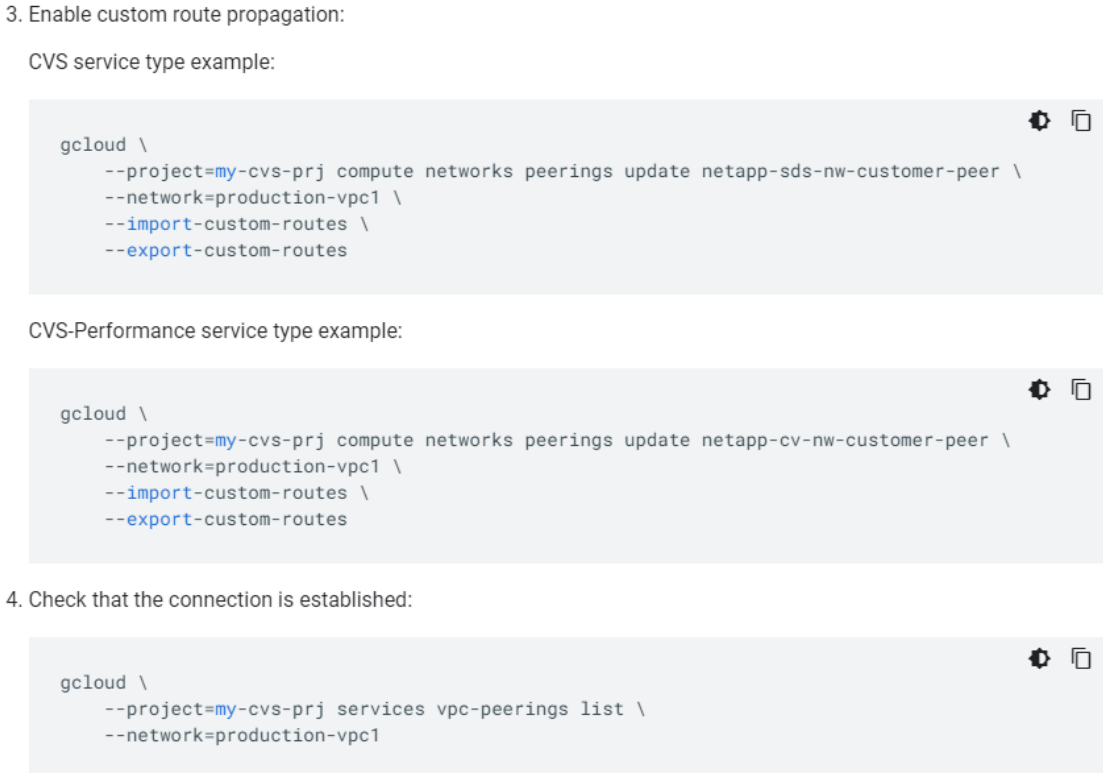
	3. Enable custom route propagation:
	CVS service type example:
	`gcloud \`
	`--project=my-cvs-prj compute networks peerings update netapp-sds-nw-customer-peer \`
	`--network=production-vpc1 \`
	`--import-custom-routes \`
	`--export-custom-routes`
	CVS-Performance service type example:
	`gcloud \`
	`--project=my-cvs-prj compute networks peerings update netapp-cv-nw-customer-peer \`
	`--network=production-vpc1 \`
	`--import-custom-routes \`
	`--export-custom-routes`
	4. Check that the connection is established:
	`gcloud \`
	`--project=my-cvs-prj services vpc-peerings list \`
	`--network=production-vpc1`
	Source: https://cloud.google.com/architecture/partners/netapp-cloud-volumes/setting-up-private-services-access